THE FUTURE OF TRUCKING, TODAY

Plus Safety Report

October 2020

Plus

Letter from the **Founders of Plus**

Plus was founded in 2016 to revolutionize commercial transport with self-driving trucks in order to create a safer and greener world. We believe automated driving technology has the potential to save lives, reduce costs, reduce carbon emissions, and power economic growth. As we have all seen during this pandemic, trucking is the engine that moves everything we use and consume in our daily lives. Self-driving trucks will generate trillions of dollars of value, reshape industries, and usher in a new generation of innovative companies. Developing this technology safely requires a world-class team, the highest safety standards, and incredible focus and patience. We have assembled a team of experts in automotive safety, self-driving technology, artificial intelligence, robotics, cybersecurity and product development to bring the safest self-driving truck system to the world.

This Safety Report highlights key principles and guidelines underpinning our safety approach and how we build safety into everything we do. We understand that most people are not experts on heavy duty trucks, much less self-driving trucks, so this report will provide an overview of the technologies behind our self-driving system, how they are used to enable self-driving trucks, and how we are developing our self-driving system to become safer than human drivers.

This Safety Report is published as a Voluntary Safety Self Assessment document based on voluntary guidance in "Preparing for the Future of Transportation: Automated Vehicles 3.0" and "Ensuring American Leadership in Automated Vehicle Technologies: Automated Vehicles 4.0" to support transparency in safe testing and deployment of ADSs.

As we continue to drive towards the commercialization of automated trucks, we look forward to sharing more about our development and our commitment to safely bringing our self-driving system to market.

David Liu   Shawn Kerrigan   Hao Zheng   Tim Daly

SAFETY IS CORE
TO OUR COMPANY VALUES

Plus adheres to a safety-first approach to developing and applying self-driving technology. It is embedded across our corporate culture, engineering process, operations, and hardware choices.

Decades of experience inform us that safety is as much a cultural challenge as it is a technical problem. Safe development, products, and operations are driven by a safety mindset established by management and cascaded through the engineering, test, and operations teams.

OUR CORE VALUES

Safety First

Safety is core to our success. It is embedded in everything we do. This goes for everything from our design and engineering approach to the training of our operations team. This core value has also shaped our approach to deployment, where we focus on rigorously tested and released products.

Always Deliver

We believe in the importance of delivering on commitments. That means we carefully consider what we commit to, and pull together as a team to ensure we always succeed.

Collaborative & Open

We value collaboration in all aspects of our work, internally and externally. We believe the best results come from sharing a broad understanding of the system. There is always room for improvement, and everyone needs to be open to feedback regardless of their role.

Realizing the vision of SAE Level 4 automated trucks will take the collective efforts of the entire trucking ecosystem. We appreciate the opportunity to work with partners and allies, including industry groups, educational associations, state and federal government, and other stakeholders to advance self-driving technology for commercial transportation.

WE ARE BUILDING AN AUTOMATED DRIVING SYSTEM FOR HEAVY DUTY TRUCKS

Trucking is a $600 billion industry in the U.S. It is not just an economic engine; trucking is essential to our daily lives. It is responsible for moving over 70% of the goods we eat and use every day[1]. Yet trucking remains one of the most dangerous careers in the country, with accidents involving trucks resulting in nearly 5,000 fatalities and over 150,000 injuries in 2018[2]. Given these risks and the challenging lifestyle of being on the road, it is not surprising that there is a massive shortage of drivers. The average age of a long-haul driver is 46, and new drivers are not entering the field as fast as drivers are retiring[3]. There was a shortage of 60,000 truck drivers in the U.S. in 2018, and the shortage is expected to grow in the coming years as drivers retire[4].

Plus is built around the vision that self-driving technology will first be successfully deployed in trucks on highways. Self-driving trucks do not get tired or distracted. They have superhuman awareness of their surroundings. They never glance away from the road for even a moment and are always looking in all directions at once. They can precisely determine the speed and direction of vehicles up to a mile away. They can locate themselves with centimeter precision on a high-definition map. They can see in the dark, in fog, and through clutter like bushes and fences. Greater situational awareness and precise control capabilities are critical for 80 foot-long heavily loaded tractor trailers with less maneuverability and longer stopping distance than cars. We apply learning across the full fleet of trucks, so that every truck in the fleet improves based on a high collective volume of driving experience, beyond what a professional human driver would experience in many lifetimes.

"We apply learning across the full fleet of trucks, so that every truck in the fleet improves based on a high collective volume of driving experience, beyond what a professional human driver would experience in many lifetimes."

Table of Contents

01 OUR APPROACH TO DEVELOPING A SAFE AUTOMATED TRUCK

We Are Taking Two Steps To Deploy A Safe Self-Driving Truck Product:

01 | Develop a fully functional and safe self-driving system

Create an automated driving system that can reliably drive a truck from a hub to a highway, down the highway, and to a hub on the other side. This system needs to handle these drives significantly more safely than a human driver.

02 | Prove it is safe enough to release without a driver

Not only does the technology need to work safely, we need to prove that it is safer than a human-driven truck before we deploy it without a driver.

Most of this report is dedicated to explaining our approach to the first step. Developing a fully functional and safe self-driving system for a Class 8 truck is a complex undertaking. Because machines ultimately cannot learn and adapt on the fly to novel situations like a human driver can, how we approach developing our safe driving system is very different from how human drivers learn to drive. Fortunately, many proven engineering techniques apply in this area, and there exist a number of standards and guidelines that we can apply in developing safe automated driving systems, such as ISO26262, ISO/PAS 21448, and UL4600. Applying these standards to our component-level systems and our overall processes will allow us to incrementally build a reliable self-driving system.

As described in the second step, before we take the driver out of the vehicle, we need to prove that our system is safer than a human driver. We believe this will require billions of real-world miles. Given the stakes involved with the size of a Class 8 truck, we don't think it would be responsible to conduct demos without a driver before accumulating the required miles and extensive test cases needed to demonstrate safety.

02 HOW AND WHERE THE PLUS SELF-DRIVING SYSTEM WORKS

IMPORTANT FOUNDATION:
A STRONG PERCEPTION SYSTEM

Progress in autonomous driving has closely tracked progress in sensing and perception algorithms. 3D lidar was the key enabler for the Darpa Urban Challenge. Deep learning made autonomous driving using a camera-only sensor suite possible for the first time in the 2016 timeframe. The new imaging radar systems make it possible to drive using only radar.

Perception is also the hardest part of the system to test. The job of the perception system is to create a structured, semantic representation of the complexities of the real world. However, in an open-world robotics problem like autonomous driving, nearly anything can happen. Since real-world phenomena are drawn from a heavy-tailed distribution, the number of test samples needed is enormous.

For this reason, we focus on sensor fusion as a foundational element of our perception system. Across all sensors that we use, we strive to maintain high detection, tracking, and localization performance for each sensing modality on its own. That provides us with redundancy, which we can use to make the perception system more reliable. During our development process, we ensure that the system is theoretically capable of driving with only a single sensing modality as a redundancy measure. However, on public roads, we always combine all three major sensing modalities to ensure the highest level of safety.

Roof sensor beam on experimental driverless prototype.

Object and Event Detection and Response

Our self-driving system is equipped with a perception system that detects objects and events, and it then communicates information on the vehicle's surroundings to our planner so that it can best navigate the road. The perception system uses a variety of sensors to give the truck a 360° view of its surroundings. These currently include lidar, radar, and cameras. Each has strengths and weaknesses, creating a perception system where they complement each other. As new sensors are developed or existing sensor types are improved, we rapidly test and integrate them in order to make our system more robust. Within each type of sensor we add to our systems, we procure best-in-class capabilities.

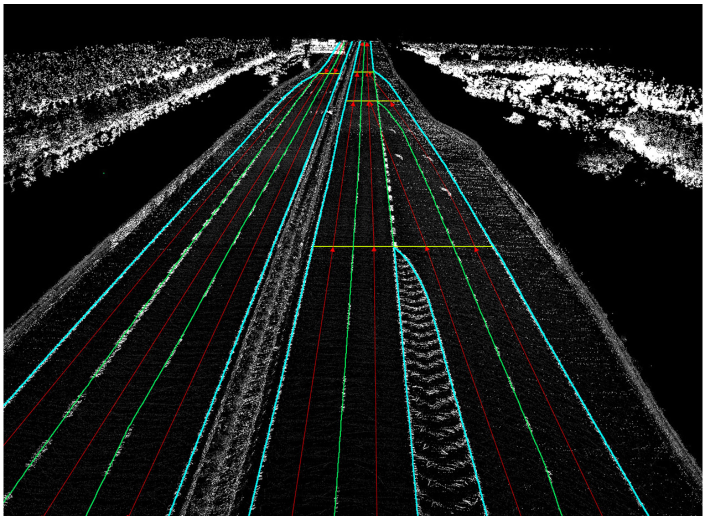

Our in-house high-definition mapping and map annotation capabilities help us improve safety by telling the perception system what it should expect to see.

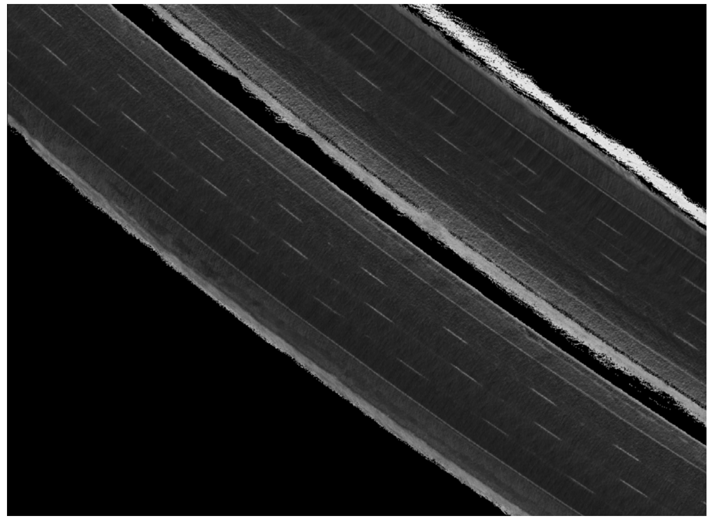

Perfectly calibrated beam-to-beam brightness makes even worn lane markings clearly visible in a lidar map.

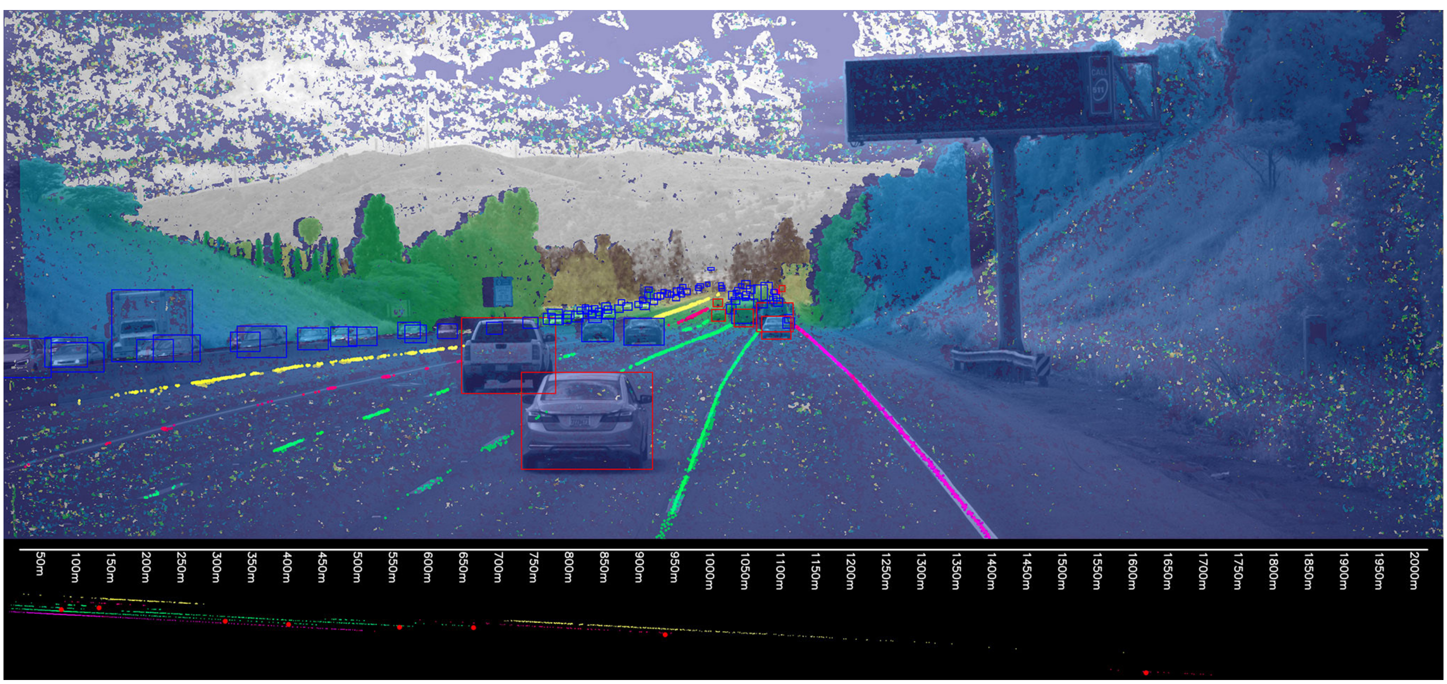

Long range perception color coded by distance to the scene. Examples of distances to some of the objects shown in red boxes are shown in the bottom part of the image as red circles.

Sensors We Use

Lidar

Lidar sensors use lasers to precisely measure 3D points. A single scan can deliver half a million points or more in just a tenth of a second, creating an accurate, though sparse, representation of all nearby objects. Using both our cameras and our lidar, we can localize with respect to our HD maps which contain structured priors that can aid in rapidly and accurately understanding our surroundings. Lidar provides precise distance information, but limited shape information and very limited texture information. A lidar will tell you where something is, but not often what it is.

Radar

Radar sensors transmit radio waves. Once those waves hit an object, they return to the sensor, which makes it possible to directly measure a range rate. This improves the accuracy of velocity estimation in our tracker.

Cameras

Cameras provide very rich data inclusive of color. Cameras have interchangeable lenses, and are passive sensors, so they can be made to see very far. However, distances to those objects need to be calculated to know the exact location. It is also more difficult to detect objects in low visibility conditions, like fog, rain, or at night.

We also use thermal cameras which allow us to reliably detect vulnerable road users even in extreme glare or darkness, clutter, and fog.

INS

Inertial Navigation Systems (INS) are used to determine the truck's position, orientation, and speed. They use the Global Navigation Satellite System (GNSS) as well as an onboard Inertial Measurement Unit (IMU) to calculate relative movement. We augment the INS using our proprietary visual Simultaneous Location And Mapping (SLAM) system and lidar localization for greater reliability and precision.

Crashworthiness

The Plus self-driving system is integrated into vehicles from truck manufacturers that meet or exceed all applicable federal safety standards. We work closely with our partners to ensure that our installation of sensors, wiring, and compute will not adversely affect any safety functions of in-vehicle systems.

The Plus sensor mounting system is independent of the vehicle structure and does not affect the base vehicle structural integrity or crashworthiness. Sensor mounting locations are selected to be incorporated into existing external features such as bumpers and mirrors or placed high above the level of possible contact with pedestrians or other vulnerable road users.

Important Role of Cameras

Cameras are the most valuable sensors for autonomous driving on highways. They can be equipped with lenses that allow them to see fine details at extreme ranges. They provide the richest semantic information about the world. They are the only sensor that can tell us the color of a traffic light and are by far the best sensor for reading lane markings and signs. Because they provide high resolution, dense information, and rich texture, their effective range can be much greater than that of lidar.

The biggest challenge for camera-based perception and the reason why camera-based autonomous driving did not become practical until recently, is that it is very hard for a computer to make sense of the images. Humans can do it very well, but we don't actually know how our vision system works. It was not until modern deep learning methods were invented that we were able to get human-level and sometimes even superhuman performance for camera-based perception. However, deep learning brings an entirely new challenge. How can we prove that such an enormous model is safe and robust in all cases? This challenge is what underlies our decision at Plus to generate massive test datasets by driving billions of miles with our perception system in a human-centered automated truck.

OUR FOCUS: SELF-DRIVING SYSTEM FOR THE MIDDLE MILE

Plus is focused on automating the middle mile of long-haul trucking (i.e., the long-haul portion of the journey on interstate highways), which accounts for around 80% of long-haul trips. Highway driving is a great fit for what cutting edge automated driving systems can currently handle. The routes are fairly repetitive, and the scenarios, while challenging, are not as complex as dense urban traffic that involves pedestrians and other forms of micro mobility.

We are training our driving system to handle commercial roads and highways in order for our self-driving trucks to go from one distribution hub onto the highway, down the highway, and then to another distribution hub on the other end. This clear and restricted operational design domain (ODD) has focused our development and paved the way for self-driving trucks to be one of the first automated vehicles to become commercially viable at scale.

PLUS IS CONTINUALLY EXPANDING OUR OPERATIONAL DESIGN DOMAIN

The Plus strategy is to know one's limits, be prepared for any situation, and have a safe response plan. It is not just a question of what to do in a critical situation, but how to avoid getting into a critical situation in the first place. Studies show that practically all road vehicle incidents could have been avoided if better decisions were made prior to the critical event. This is not just the benefit of hindsight, but also being totally aware of the scenario and planning for all eventualities.

The automated vehicle must be designed and demonstrated to have a sufficiently high safety performance in the intended areas of operation. It must be expected that there will be operational conditions which are beyond the capability of the self-driving vehicle, or a human driver, for that matter. The key to the safety of self-driving vehicles is to understand and avoid such operational conditions.

Plus applies a rigorous methodology to managing the ODD concept. The ODD is the set of geographic regions, environmental factors, surrounding traffic scenarios, and required maneuvers in which the self-driving system has very high confidence of safe execution. The challenge is to increase the depth and breadth of the ODD. This involves both improving performance within a defined region and additionally increasing the extent of the region, conditions, and scenarios of capability. At Plus, ODD is managed at various levels within the self-driving software stack. Each component in our system can independently signal that it is going out of, or has gone out of its ODD. Our watchdog makes the call as to which components are critical. It will transition the truck's ODD status as appropriate.

"The ODD is the set of geographic regions, environmental factors, surrounding traffic scenarios and required maneuvers in which the self-driving system has very high confidence of safe execution. The challenge is to increase the depth and breadth of the ODD."

ODD Checker

The operation of the automated system out of the defined ODD could lead to unsatisfactory performance. An undesirable situation could arise if the autonomous system does not recognize that the current or future scenario is not within its ODD. Consequently, Plus has adopted the concept of the ODD Checker. Every Plus feature includes a structured ODD definition and a set of tests for determining whether the operating conditions meet the defined ODD. This emphasis on monitoring for a valid ODD extends to predicting and reacting to future eventualities.

"Plus has adopted the concept of the ODD Checker. Every Plus feature includes a structured ODD definition and a set of tests for determining whether the operating conditions meet the defined ODD. This emphasis on monitoring for a valid ODD extends to predicting and reacting to future eventualities."

The combinations of the ODD scope and the ODD checker result in four possible states. Clearly operation in state C, which is an unknown suitable operation region, should be minimized. State D, which is an unknown unsuitable operation region, must be avoided.

Subsystem		ODD Scope	
		Outside	Within
ODD Checker	Valid checker condition (known ODD status)	**B** Outside ODD and we know it	**A** Within ODD and we know it
	Invalid checker condition (unknown ODD status)	**D** Outside ODD but we don't know it	**C** Within ODD but we don't know it

The following steps are used to establish confidence in the operation of an automated truck:

01 Define the ODD Scope - the region and conditions for safe and reliable automated operation.

02 Develop an ODD Checker - a method to detect when our vehicle is within or outside the ODD.

03 Test and validate the ODD and ODD Checker.

04 Use the ODD Checker as part of the self-driving Auto Mode enablement criteria.

05 Use the ODD Checker to detect when the vehicle is out of or about to go out of ODD and the put system into a safe state.

Once there is a determination that the ODD conditions may be violated, now or in the immediate future, the system will transition to a safe state. The safe state may be as subtle as switching to a redundant system, transfer to manual driving (for a driver-in system), or a reduction in performance such as lowering speed or a safe stop condition. The system response will be scenario and system state dependent.

Minimal Risk Conditions

Despite the care exercised in design, any system can encounter unexpected failures in sensing, interfacing, or software. Moreover, changes in external conditions or driving maneuvers could cause the system to move outside its ODD. The system must first and foremost be capable of detecting the event, following which it must react appropriately to enter a minimal risk maneuver in response.

Our self-driving system includes supervisory modules that consistently monitor the health and activity of sensors, communication interfaces, power distribution, computing infrastructure, and the software executing on it. These modules also monitor changes in the operating environment that could cause our trucks to be outside their ODD. When the system detects such a failure or event, depending on the type of event and the potential sensing and actuation options still available, the system executes an appropriate minimal risk maneuver. Minimal risk maneuvers include pulling off the highway at the next exit, slowing and pulling to the side of the road, stopping in place, or other maneuvers that bring the truck to a minimum risk condition for our vehicle and other road users. We understand that stopping in an active roadway can create a safety hazard, so a minimal risk maneuver would only be done as a last resort.

Data Recording

Our trucks have been designed from the ground up to record all sensor data any time the vehicle is on. We can replay this data to evaluate improvements to our software on an ongoing basis. The recorded data includes:

- Vehicle state, including whether it was in autonomous or manual mode
- Location data
- Sensor data (images from cameras, point clouds from lidar, and detection data from radar)
- Outputs and intermediates from our key algorithms and models in the perception, prediction, localization, planning, and control components
- Footage from an interior camera that captures driver actions.

Our sensors are synchronized to a highly accurate clock source, and all intermediate processing results as well as control outputs are also timestamped using this clock source. That enables our replay simulation system to precisely and deterministically recreate everything that happens on our trucks. While the fundamental architectural design choice to record everything at such high fidelity imposes a great burden of data upon our back-end systems, we believe it is important for the development of metrics to ensure our system is continuously improving, providing notable scenarios to advance system learning rapidly and aid in event reconstruction.

Post-Crash ADS Behavior

If a Plus vehicle is in a crash, the vehicle would immediately notify a command center of the event as well as transfer some basic data for review if the connection is not affected by the crash event. If the crash event affects communications, the missing heartbeat from the vehicle is used by the command center to initiate protocols to determine the vehicle location and state.

The watchdog evaluates the status of sensors, communication interfaces, compute unit, power distribution, and software and executes a minimal risk condition maneuver that aligns with the current state of the vehicle and the external scenario that can be determined reliably. The vehicle and/or operation center will invoke services necessary to address the situation. Depending on the location and the severity of the event, the services may be first responders or company personnel.

The autonomous vehicle will also carry documentation that allows someone to quickly check the status of the vehicle, as well as operate the vehicle manually to move it out of the roadway if necessary. Out of an abundance of caution, for any crashes, the vehicle would disable autonomous driving until evaluated by qualified personnel and deemed okay to engage. The credentials of personnel authorized to make this decision would be controlled by hardware authentication devices. The vehicle is labeled with a 24/7 phone number where the Plus operation center can be reached so that anyone dealing with the vehicle post-event can be clearly instructed on what to do in addition to on-board directions.

Plus often has several vehicles on the roads simultaneously for development and testing. Depending on analysis of early data from a particular crash event, Plus may decide to suspend autonomous driving across the entire fleet if warranted. This decision may involve different types of responses based on the details of the event and probable root cause.

03 BUILDING SAFETY INTO OUR SELF-DRIVING SYSTEM

SAFETY IN OUR SYSTEM ARCHITECTURE

Safety at Plus does not exist as a separate feature, but is built into the architecture, processes, and operations at all levels. The Plus system architecture is designed around redundancy and fail-operational fault management. Two or more levels of software or hardware exist to perform or check the operation of each single function. When a fault is detected, the system will continue operation at a safe operating level, possibly compromising performance or efficiency, but never safety.

The Plus perception stack uses multimodal data fusion. Inputs from a redundant suite of complementary technologies, cameras, lidar, and radar are fused together to create a multidimensional view of the world. The system is designed to be tolerant to faults (obstruction or failure) of any single sensor or sensor group.

Redundancy also extends to the software and compute platform and down to the interfaces to the base vehicle. State of the art computation, steering, and braking systems are sourced from the world's most competent component suppliers. We are always evaluating alternative systems to find the best solution in each domain.

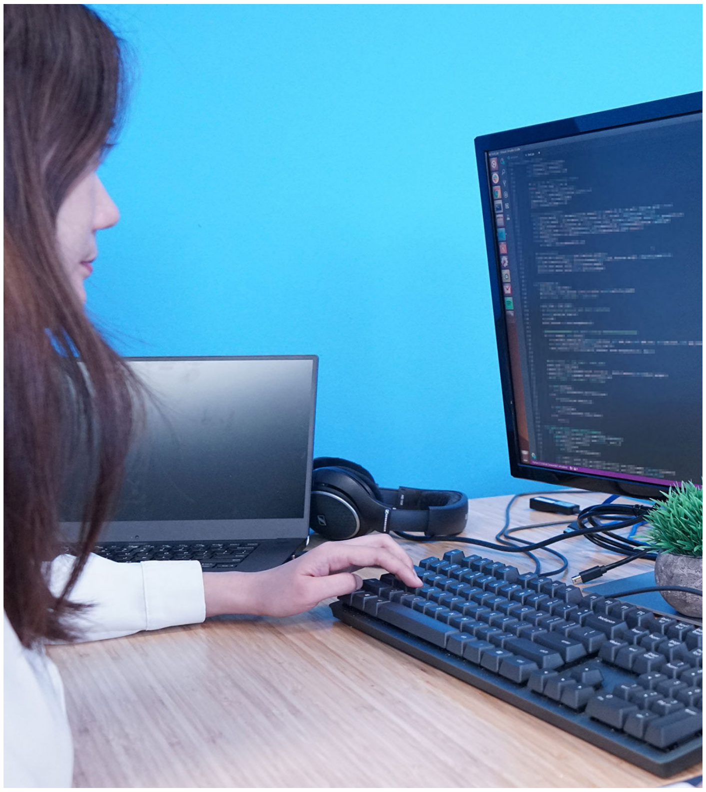

SELF-DRIVING VEHICLE CYBERSECURITY

Cybersecurity for self-driving trucks starts with physical security. This involves making sure that there is no unauthorized access to the vehicle or its systems, starting with policies and training of operations staff. Moreover, we use defensive measures to deter and detect physical access to the system.

The next aspect of cybersecurity is software security. We follow the standard software security practice in mission critical systems to ensure the integrity of the system.

Network security is also crucial to cybersecurity. For accurate location determination, fleet management, software updates, and other reasons, all vehicles must maintain regular, but not necessarily continuous, internet connectivity. We use firewalls to block incoming access, limiting any connectivity to being originated from within the vehicle to the predetermined set of services. Additionally, the self-driving software is run locally and not dependent on any remote real-time commands. Thus, malicious actors are unable to intercept or spoof any communication critical to continued operation of a vehicle.

Physical Security

- Policies and training
- Hardware authentication devices

Software Security

- Signed software updates
- Certificates

Network Security

- Firewalls
- Private network

SAFETY AT EVERY LAYER

In addition to the safety checks and redundancies that we build into our system architecture, each layer in our software stack has its own safety validation process. This is important because while each layer in our system takes some responsibility for the overall correctness, it defers some to the next layer up. For example, if the perception system misleads the planner about the state of the world around the truck, the planner may have no way to detect and correct the problem. By validating at each layer, we ensure everything is safe.

In our first layer are perception and localization. Perception describes the way our software processes and makes sense of all data about the truck's surroundings collected through our sensors. Terabytes of data are aggregated for our system to learn over time how to correctly and consistently identify an object, where it is, its speed as well as determining the location of our truck. We use statistical validation to show our confidence that what the system says it detects is in fact true. Perception, in particular, is among the most challenging areas of self-driving because of the numerous permutations and anomalies that can arise in the real world. This is also why we believe that the statistical validation of a self-driving system to be safer than a human driver will require billions of miles of public road testing. This ensures that anomalies or edge cases have become not so abnormal to the driving system over time given its repeated exposure to the real world.

Localization informs the autonomous vehicle where it is on a map and its position relative to static elements of the scene including things like lanes, road boundaries, traffic lights, and stop lines. Our localization module can make use of a High Definition (HD) map, when available, and using a combination of GPS, IMU, lidar, cameras and radar positions the autonomous vehicle to an accuracy of a few centimeters. Even in the absence of an HD map or when we detect a discrepancy between the HD map and our online sensor data, our system can make use of standard definition navigation (ADAS) or semantic maps to drive and navigate safely. This latter capability is made possible based on the rich understanding of the road semantics and geometry provided by the perception system. Indeed, our experience suggests that a vast majority of highway driving is possible without the need for HD maps, relying only on ADAS maps and perception. As with the perception system, localization requires validation against a sufficiently large set of real-world data in a variety of traffic and environmental conditions.

SOFTWARE LAYER	TYPE OF SAFETY VALIDATION USED
Perception & Localization	**Statistical Validation** We use statistical validation to prove that what the system says it detects is in fact true.
Planning & Control	**Mathematical Validation** Formal mathematical models help to prove or disprove the accuracy of what our system derived it should do next.
Drive-By-Wire & Actuation	**Functional Safety** Functional safety processes validate that the action the system control is about to take is in fact safe based on its operating design domain.

After perception and localization come prediction, planning, and control. Once the object is accurately detected, and given the truck knows its own location, the prediction system will help our planner understand the likely behavior of other road users. Then, the system will plan what the truck should do next, incorporating an assessment of the intention of other vehicles, and it will control the vehicle to make such a move. In this layer, we use formal mathematical models to prove or disprove the accuracy of what our system derived it should do next. For example, the system needs to determine how close it can get to the vehicle ahead and still be at a safe distance. The mathematical model will show whether that distance calculated by the system is indeed safe. The control function collaborates with planning to ensure that the truck does not do anything unsafe and keeps the vehicle close to the requested trajectory from the planner. If it cannot do that, it must signal out-of-ODD.

The last layer is the actuation and drive-by-wire system. Drive-by-wire makes it possible to control the vehicular systems including the acceleration, braking, steering and other functions of the truck despite no physical or mechanical linkage of the different areas. Safety validation in this layer is tried and tested through decades of developing traditional vehicles.

Functional Safety processes are being tailored and extended for autonomous vehicle application. ISO26262 Road Vehicles - Functional Safety aims to address possible hazards caused by the malfunctioning behavior of electronic and electrical systems in passenger cars. It is considered best practice but is not a regulatory requirement. It also doesn't apply to many critical functions in the autonomous software stack that use machine learning.

At Plus, it serves as a basis for a more comprehensive safety strategy. There are other standards for best practice under development including ISO/PAS 21448 Safety of the Intended Function (SOTIF) and UL 4600 Standard for Safety for the Evaluation of Autonomous Products which we are closely monitoring for their applicability.

Functional Safety Process Workflow

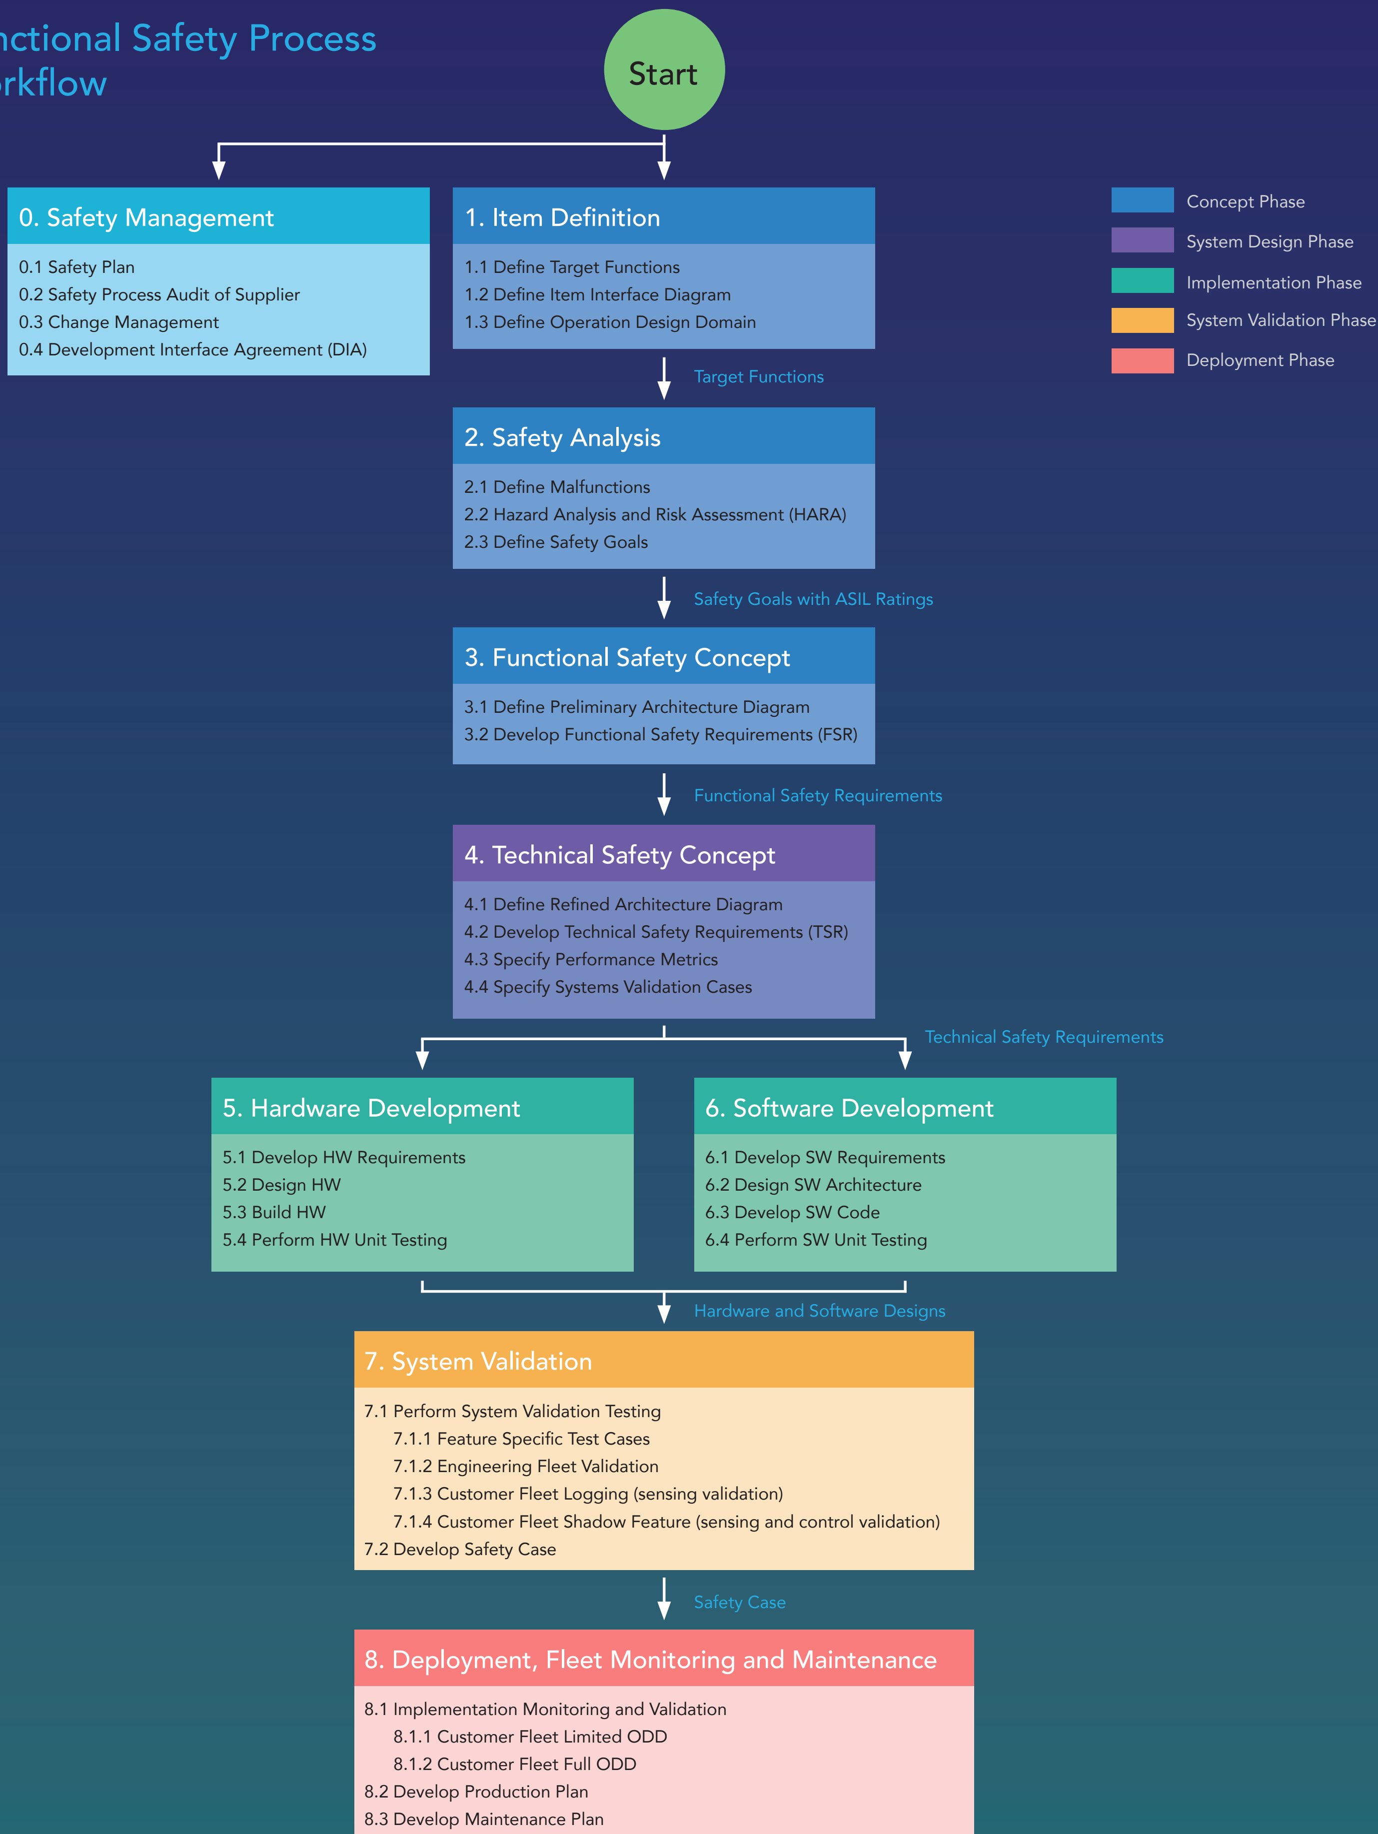

SAFETY IN OUR ENGINEERING DEVELOPMENT PROCESSES

Structured engineering development processes are employed to define requirements, provide traceability, define test and validation, and manage issue resolution.

The implementation of any safety strategy is built upon good engineering through a structured engineering development process. The framework for the Plus development process is outlined in the diagram below. The process involves a series of development stages with deliverables defined by the workflow highlighted on the outside of the development 'V'. The tools and implementation in the inside of the 'V' outline the methods for managing and recording the work products.

The Safety Concept is considered at the start of the development process and the safety goals drive requirements alongside performance and convenience features. Requirements are cascaded through the process with every defined item having a measurable performance and test method. The requirements and test cases are transferred to the test and issue tracking tool for implementation and test execution. Finally the test results are transferred to the requirements tool for coverage analysis and reporting.

Plus employs a multi-tiered test and validation program. Software features are extensively peer reviewed, tested in a comprehensive simulation environment, strategically tested on closed test tracks and only then deployed to public roads under the strict supervision of a safety driver and vehicle operations specialist. Hardware and systems are similarly tested at various levels. Plus also believes in the separation of responsibility between the development engineers and the test process. Requirements definition, system level test procedures and validation plans are managed independently of the development teams.

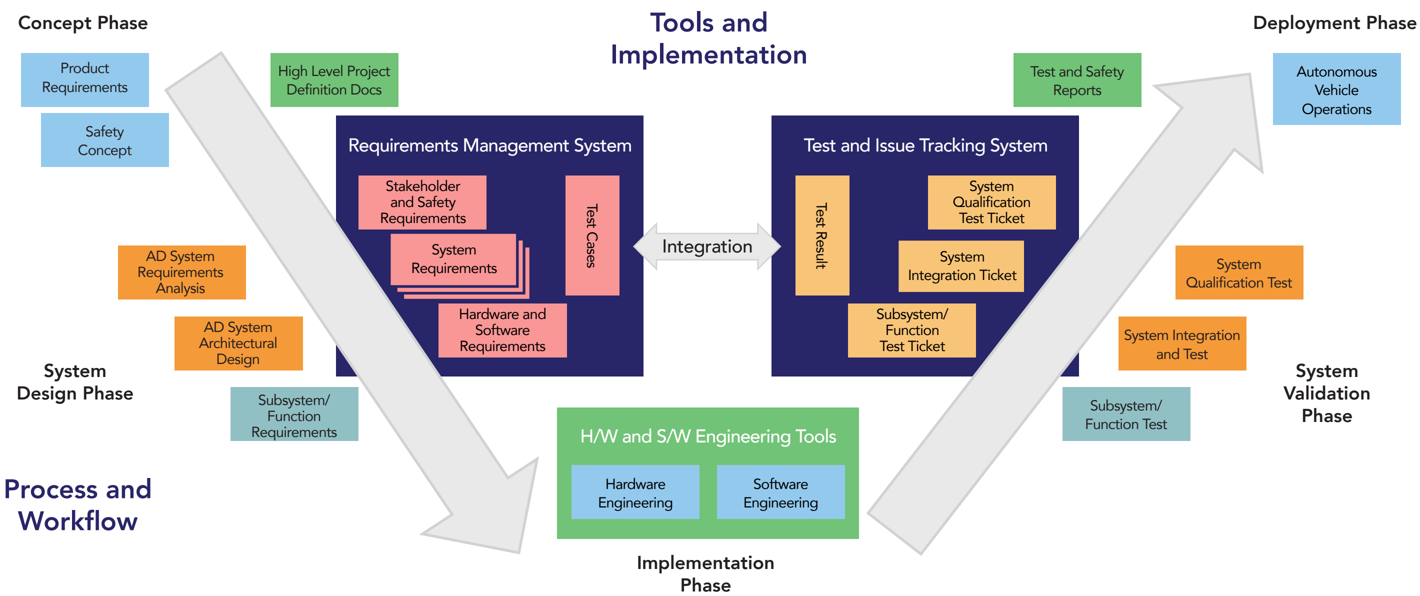

> **"Plus also believes in the separation of responsibility between the development engineers and the test process. Requirements definition, system level test procedures and validation plans are managed independently of the development teams."**

SAFETY IN OUR DAILY OPERATIONS

As the primary group interacting with our trucks, safety is top of mind for the Plus operations team. We have assembled a team of highly qualified safety drivers (seasoned Class A drivers with extensive experience operating Class 8 tractor trailers) and vehicle operations specialists (former trainers at top autonomous car companies). Each vehicle is always staffed with a safety driver and a vehicle operations specialist to ensure checks and balances in our safety processes.

The safety driver is responsible for performing a standard commercial pre-trip inspection of the vehicle before the start of testing and also ensures safe operation of the vehicle while on the road. The vehicle operations specialist conducts a separate pre-trip inspection focused on the software. They also certify that the safety driver is driving safely and has knowledge of potential observed threats, oversee the vehicle's software while in operation, and identify opportunities for data enrichment throughout the trip.

Both roles require the passing of a multi-step interview process where the candidates are evaluated for their commitment to safety and ability to recognize potential risk scenarios.

The Interview Process Involves:

1. An in-depth background and driving record check
2. A driving and risk-recognition skills assessment
3. A sit-down behavioral and experiential interview with members across multiple teams at Plus.

After passing the interview process, operations team members undergo a specialized 3-week training program which includes in-classroom and in-vehicle instruction before they are cleared to operate the vehicles.

Our drivers and specialists understand the critical nature of safety to our operations and their role in the execution of safe practices. They are trained to look out for and react to potential threats and are instructed to disengage our self-driving system and take control of the vehicle whenever necessary. The team is safety-conscious in their daily operations by verifying the vehicle and system health, prequalifying planned routes to account for commercial and autonomous vehicle compatibility, and maintaining an open line of communication with engineers to update them on their findings and experiences during test runs.

04 TEST AND VALIDATION METHODS

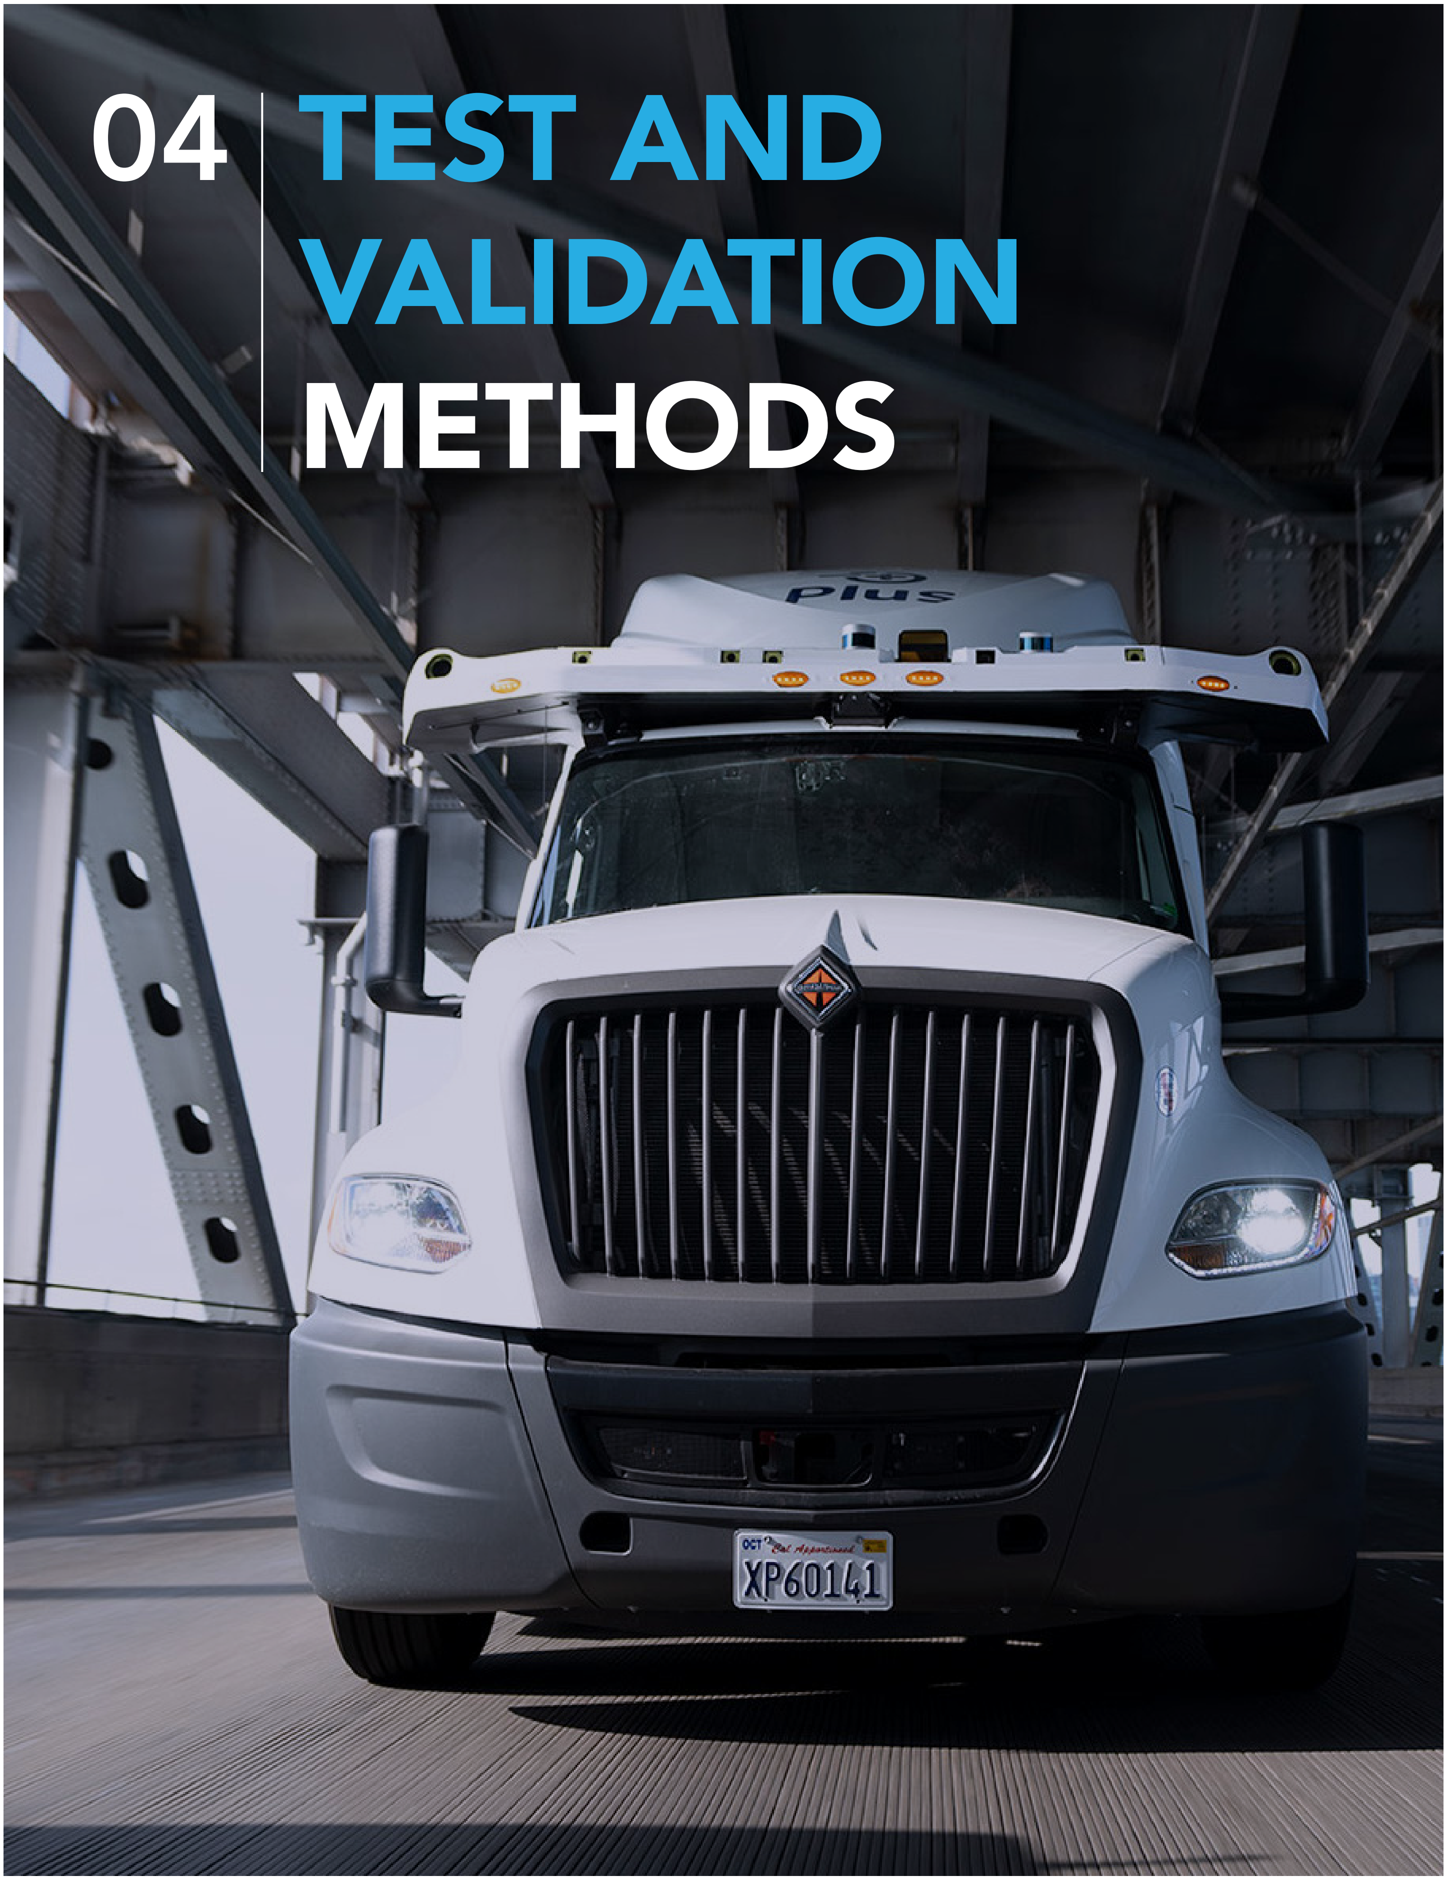

Test and validation are central to developing an autonomous system for the real world. We use a tiered validation strategy which starts at the software verification stage, transitions to extensive simulation, then software and hardware-in-the-loop (SIL and HIL) test (real software and hardware tested under simulated conditions), before proceeding to an actual vehicle. The vehicle validation starts with closed course testing before the vehicle is approved for public road testing. Once on the road, the vehicle operation is confined to a limited ODD which gradually expands to the target ODD as safety and capability is proven. Each test and validation stage is important and designed to confirm the effectiveness of a new feature or code fix, and allows us to continually refine and enhance our system. We carefully select sensors, harnesses, actuation systems, power distribution networks, communication interfaces, computation systems, and redundant vehicle foundational systems subject to automotive grade performance standards for temperature, vibration, humidity, dust resistance, and mechanical shock, among other criteria. In addition, we test vehicles at a system level to catch any requirements that may fall through the cracks in integration. Testing for hardware robustness includes real-world testing in extreme conditions as well as exposure to simulated conditions for accelerated testing. For instance, vibration and shock testing include both rumble tracks and accumulated high mileage runs on actual highways.

"We use a tiered validation strategy which starts at the software verification stage, transitions to extensive simulation, then software and hardware-in-the-loop (SIL and HIL) test (real software and hardware tested under simulated conditions) before proceeding to an actual vehicle."

We test and validate our software via a combination of miles and scenarios. More miles ensures that we see a larger variety of scenarios while simulated scenarios allow us to reach an entirely different scale and test thousands of randomized scenarios before each release goes on the road. We augment the simulation and replay testing process with a highly streamlined repetitive testing process facilitated by our operations team. Each release that passes the above steps is released to the operations team with copious notes on what changes might affect them and then fed into a sequence of real-world tests that scales up from closed track testing to low speed road testing to highway testing. We brief the safety drivers and vehicle operations specialists at each release on what to expect. They in turn evaluate the software's performance and participate in reporting meetings with representatives of the development team to share observations in person.

In addition to in-person meetings and manual reports, we rely on statistics from recordings to evolve and maintain metrics of performance such as number of disengagements to ensure that we are constantly improving performance.

SIMULATION

Simulations are used to ensure that our automated driving system is always improving. All of the driving experience we are accumulating, as well as synthetic scenarios that we create, are reproduced to check each new version of the system before it hits the road. Here we introduce the four main simulation platforms that we use.

01

Scenario Simulation

The scenario simulation is a closed loop test which is designed to test the software in specific scenarios. We use two ways to generate test scenarios.

To create “synthetic scenarios”, our experienced testing engineers use our powerful in-house GUI scenario editor. This allows us to define the various surrounding driving vehicles and their behaviors in those scenarios. We define our scenarios based on ISO/WD 34502 (Road vehicles — Engineering framework and process of scenario-based safety evaluation), SOTIF (ISO 21448) and other regulations or standards, in order to build a complete scenario library.

On the other hand, “log-based scenarios” are automatically generated from the data we have mined from our public road tests. These have been tagged in our event miners as interesting scenarios (e.g. heavy-traffic, sudden cut-in, etc.).

Each scenario simulation yields a few types of results. These include different quantitative metrics and an overall pass/fail evaluated against the expectations. The evaluation tool also allows us to compare two or more evaluation runs to see which one resulted in better metrics. Finally, a simulation video is available so that our team can view in detail to ensure the desired results were achieved.

Our software is run against these synthetic and log-based scenarios periodically and on-demand, to ensure that a new version of our software will not regress on existing tests and incremental improvements can be achieved.

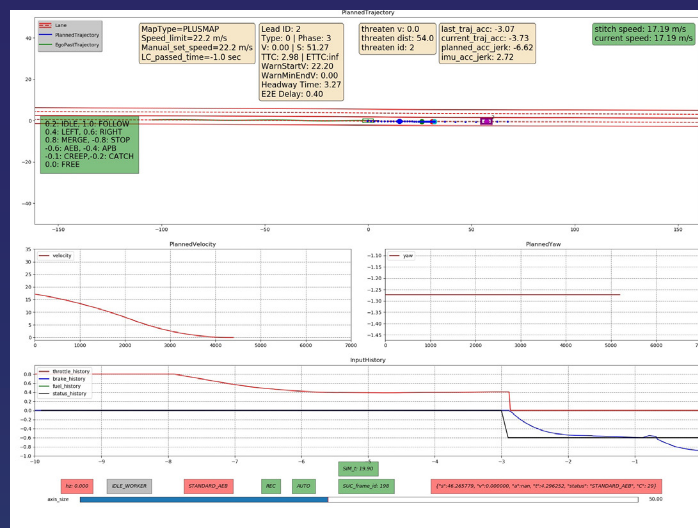

A planner diagnostic dashboard from scenario simulation.

02

Large-scale Replay Simulation

The large-scale replay simulation is designed to test software in a way that is closest to public road testing. We replay the data collected from our operating routes, and run our full software stack (from perception to control) on it to get the full mileage simulation result.

To do this, we build a very precise vehicle dynamics model to simulate the truck with different trailers first. And we also build proprietary technology to allow us to tweak the surrounding driving vehicles and road geometry based on ego truck driving state. This replay simulation platform enables us to discover potential problems on our new software before running our truck on public roads.

03

Perception Simulation

The perception simulation is designed to test our perception algorithm under different scenarios and conditions. For example, in lane detection, we may change the scenario to have the truck drive a little more left than the actual driving to check the performance of the lane detection algorithm.

04

Vehicle-in-loop Simulation

The vehicle-in-loop simulation (ViL) is designed to test the system on real trucks, under simulated scenarios. This is usually done on a closed course.

Take the testing of emergency braking for example. We deploy the new software on the truck first and then drive the truck in a wide area. We could simulate a fake dynamic vehicle which cuts into our lane in a very close leading distance. That way we can check the real braking reaction time of our truck, which is safer to do than public road testing.

Sensor layout studies and auto-calibration ground truth generation in the CARLA simulator.

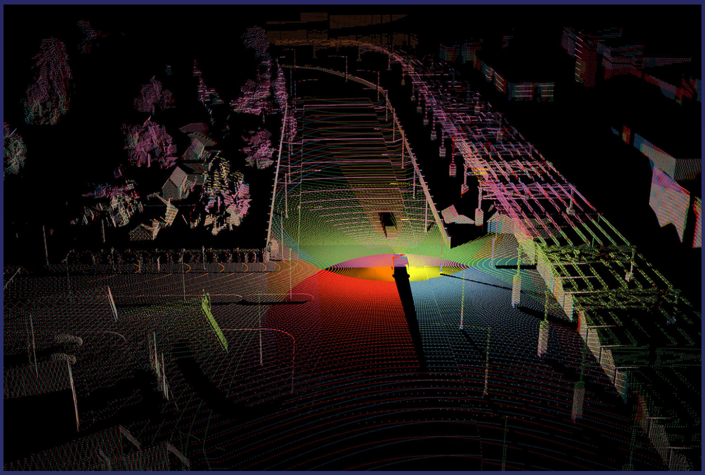

CLOSED COURSE TESTING

Closed course testing is particularly important for scenarios that are otherwise unsafe to test on a public road. One example is the ability of our driving system to detect and automatically engage the emergency brake when a pedestrian is detected in front of our vehicle on the highway. We have also partnered with the Minnesota Department of Transportation (MnDOT) to conduct cold weather testing in snow and ice conditions on their MnROAD closed course since these conditions are not native to the California Bay Area where we are headquartered. Plus has conducted closed course testing on multiple tracks throughout the U.S.

PUBLIC ROAD TESTING

Public road testing is invaluable and the only way to truly show that our self-driving system is safe and works in the real world. Simulating the real world well enough that it can replace road testing is harder than building a self driving vehicle. The fact that the system works in the simulator does not provide strong statistical evidence in the case of truly unexpected events -- landslide, fire, hurricane, tampering and vandalism, malicious drivers, mechanical failures on ego vehicle, or new obstacle types.

We employ a clear and stringent process to ensure that public road testing is done in the safest manner. Engineers first fill out detailed test requests with required information about the testing desired. A senior member of the operations team then reviews these test requests, attains any necessary clarifications from the test requestor, and then schedules the test for execution with an operations team consisting of one of our experienced Class A drivers and a vehicle operations specialist. The assigned operations team also has the opportunity to ask additional clarification questions before beginning the test.

Afterwards, the onboard vehicle operations specialist will provide detailed in-vehicle and post-run notes that are communicated to the test requestor. For items discovered during the test that may need extended engineering attention, vehicle operations specialists will submit this information into our testing management system. This ticket is then assigned to a dedicated member of the engineering team and addressed in a timeframe suitable for the priority level. Based on this feedback and the resolution of associated tickets, test requestors may request another iteration of the road test.

We understand that public road testing comes with the highest degree of responsibility. We do not take this responsibility lightly, and only conduct public road testing for software changes that we deem ready and that are truly necessary.

The Real World Is One Complex Place

For all the scenarios our computers can simulate, there's no substitute for the complexities we see in the real world, especially the rare cases we encounter. Our self-driving system can only be deemed safe enough to not require a safety driver in the vehicle when these seemingly extraordinary events today can be handled as ordinary events.

Example: Turbine blade followed by chase vehicle with flashing lights.

Example: Truck carrying tree.

INDEPENDENT TESTING

Independent testing is a cornerstone of our plans for commercial deployment. We work with external testing facilities to develop and validate a range of autonomous features for performance, robustness and safety. While we do an enormous amount of tests ourselves, independent testing is free of any bias from our team. If human drivers have to pass a driving test to obtain a commercial driver's license to operate a semi-truck, we believe third party validation should be required of self-driving systems as well. We are pleased to have contracted a deeply experienced testing team with the largest testing ground in the U.S. to provide independent test planning and execution of our self-driving system, to give us and the public the assurance of our system's capabilities[5].

The development of the independent test program started in early 2020 after a thorough review of Plus features and ODD. The testing will evaluate the Plus self-driving system's ability to consistently handle scenarios that best simulate complex, realistic driving conditions, with multiple vehicles operating in the vicinity of the Plus truck. The independent test plan is driving the Plus in-house validation activities and will be repeated and expanded using the state-of-the-art facilities in Ohio. The facilities enable self-driving truck operation across various operating conditions, many of which would be considered hazardous to test on public roads.

The planned testing is the first phase of a more comprehensive and ongoing test program that will build on this experience to ensure a safe and validated solution as the Plus ODD expands to a larger proportion of the nation's highways.

HUMAN-MACHINE INTERFACE

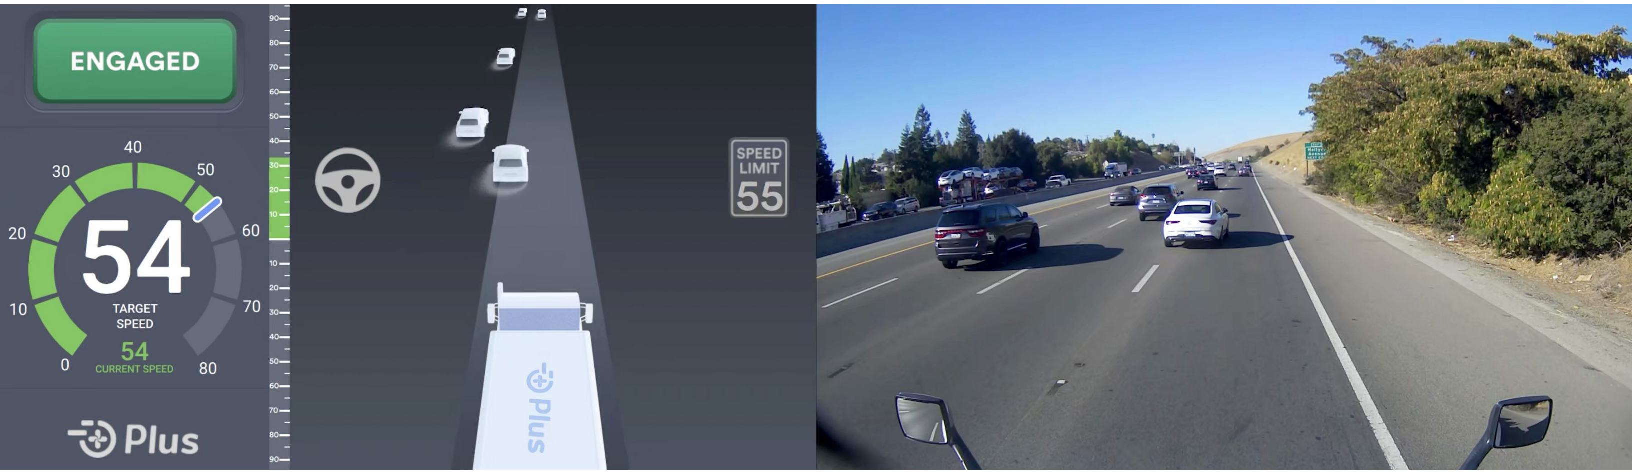

The human-machine interface (HMI) on our test vehicles is designed for our safety driver to use during our development phase. Our HMI uses audio, tactile, and visual interactions with the driver and vehicle operations specialist. The primary interface for the driver is a switch to engage / disengage autonomous mode. However, the driver can also manually intercept control at any time using the steering wheel, brake pedal, or throttle to disengage autonomous driving. Moreover, the driver can hit an emergency stop button that electrically disengages the autonomous driving system, so that the driver can have full control over the vehicle.

The vehicle operations specialist and any additional passengers in the truck (e.g., engineers and demo riders) have a full user interface displaying the system's current state; the roadway around the truck; its planned trajectory for the vehicle; its perception of the environment around it, including static and dynamic objects; and a rendering of which objects are deemed a threat in a form easy for humans to perceive at a single glance. The current state includes the driving mode of the vehicle (autonomous / manual), current and target speeds, and control inputs being applied to the vehicle by the system including steering, throttle, and brake, if in autonomous mode.

Separately, a simple display interface indicates the health of the system including sensors, communication interfaces, power distribution, computing infrastructure and software, as well as current recording status and whether the vehicle is in its ODD. This ensures that the safety driver and vehicle operations specialist are always clear on whether the system is okay to drive as well as whether it is ready to switch to autonomous mode.

Upon completion of development, the vehicle is expected to be completely automated without an vehicle operations specialist or safety driver in it, in which case, the human machine interface will include an interface between the truck and the other users of the roadway, including drivers of other vehicles, pedestrians, and other autonomous vehicles. This interface is expected to use traditional means such as turn signals, brake lights, head lights, and horn. The self-driving system design is planned to take advantage of all these means as necessary to communicate with other occupants of the roadway.

05 CONCLUSION:
THE ROAD TO COMMERCIALIZATION

Conclusion: The Road To Commercialization

We have come a long way in the development of our self-driving technology in the last several years. We are excited that we are putting our self-driving system to work daily for Fortune 500 private shippers, and continue to refine our technology to serve their shipping needs. We continue to work closely with partners across the trucking ecosystem to bring our self-driving technology to market.

As we prepare for commercialization, with mass production starting in 2021, we remain committed to safety first so that all fleets can benefit from the safety and fuel efficiency gains from deploying our self-driving system.

CONSUMER EDUCATION AND TRAINING

As developers of automated driving systems, Plus believes it has a shared responsibility to educate those who will be using and interacting with the technology on its implications. As we start expanding this group from internal to external entities (shippers, distributors, and consumers), we are dedicated to creating informative training resources to ensure all parties have the proper knowledge to support safe interactions with our self-driving system.

FEDERAL AND STATE REGULATIONS

Supportive and clear policies are key to the safe rollout of self-driving trucks. We ensure that all vehicles in our test fleet are correctly registered, insured, and licensed for normal operations. We design our self-driving trucks to safely follow all federal, state, and local laws and we intend to exceed regulatory requirements. Additionally, our vehicles are designed to meet all federal motor vehicle safety standards, working in close collaboration with truck manufacturers.

As regulations for self-driving vehicles continue to evolve, Plus is working closely with industry peers, industry associations, and regulators at local, state, and federal levels to share information and collaborate to support the safe development and deployment of automated trucks.

GOING FULLY DRIVERLESS ONLY WHEN IT IS STATISTICALLY PROVEN SAFE

We can't underscore enough that while self-driving trucks will deliver significant safety and economic benefits, the rollout of fully driverless trucks can't be rushed and must be done safely with the proper statistical validation. We believe this will take billions of real-world miles to prove safety before removing the human driver. While a daunting number, when it comes to people's lives, we must do what is right. Plus is committed to taking the safest road to the commercialization of our self-driving system.

Endnotes

[1] "ATA Freight Forecast projects 25.6% rise in tonnage by 2030." Bulk Transporter. 28 Aug. 2019, https://www.bulktransporter.com/resources/article/21658181/ata-freight-forecast-projects-256-rise-in-tonnage-by-2030. Accessed 2 Oct. 2020.

[2] National Center for Statistics and Analysis. (2020, September). Quick Facts 2018 (Research Note. Report No. DOT HS 812 951). National Highway Traffic Safety Administration.

[3] Miller, Eric. "ATA: Trucking Industry Was Short More Than 60,000 Drivers in Meeting Demand at End of 2018." Transport Topics. 24 Jul. 2019, https://www.ttnews.com/articles/ata-says-truck-driver-shortage-course-double-decade. Accessed 2 Oct. 2020.

[4] American Trucking Associations. (2019, July). Truck Driver Shortage Analysis 2019.

[5] "Plus.ai Pushes Boundaries by Using the Industry-Leading Multi-Vehicle Approach to Independently Test its Self-Driving Trucks at Transportation Research Center." Plus, 21 Jul. 2020. Press release.